

# Towards a new generation of cryptographic software

Peter Schwabe

Max Planck Institute for Security and Privacy  
Bochum, Germany

10. December 2025

# A new generation of crypto software?



# A new generation of crypto software?



- ▶ Post-quantum



- ▶ Post-quantum
- ▶ **NIST-PQC was revealing**
  - ▶ Wrong proofs
  - ▶ Hidden assumptions in proofs
  - ▶ Proofs that did not apply to the scheme
  - ▶ Software that did not implement the scheme
  - ▶ Bugs in software
  - ▶ (Compiler-introduced) timing leaks





- ▶ Post-quantum
- ▶ **NIST-PQC was revealing**
  - ▶ Wrong proofs
  - ▶ Hidden assumptions in proofs
  - ▶ Proofs that did not apply to the scheme
  - ▶ Software that did not implement the scheme
  - ▶ Bugs in software
  - ▶ (Compiler-introduced) timing leaks
- ▶ **PQC migration as a chance**
  - ▶ Applications *need* new implementations
  - ▶ Piggyback on migration to deploy better software



- ▶ Post-quantum
- ▶ **NIST-PQC was revealing**
  - ▶ Wrong proofs
  - ▶ Hidden assumptions in proofs
  - ▶ Proofs that did not apply to the scheme
  - ▶ Software that did not implement the scheme
  - ▶ Bugs in software
  - ▶ (Compiler-introduced) timing leaks
- ▶ **PQC migration as a chance**
  - ▶ Applications *need* new implementations
  - ▶ Piggyback on migration to deploy better software

**Goal: Strong link between mathematical object and implementation**



*I hope you're well!*

*"Gilles Barthe, Ben Laurie, Dan Boneh, Meredith Whittaker, and myself would like to invite you to a workshop on "High Assurance Crypto Software".*

*[...]*

*The goal: Bring together the best crypto-implementers and formal-methods verification people from open source, industry, and academia. Introduce them and their projects to each other, and scope and develop practical collaborations that improve verification of crypto code."*

—Trevor Perrin, late 2015

# Fast forward 4 years (February 2020)





Goal: Formally verified implementation of Kyber

≈9 papers, 30+ collaborators


Basavesh Ammanaghatta Shivakumar,  
**Santiago Arranz Olmos**, José Bacelar Almeida,  
Gustavo Xavier Delerue Marinho Alves, Manuel  
Barbosa, **Francisca Barros**, Gilles Barthe, Lionel  
Blatter, Chitchanok Chuengsatiansup, Ignacio  
Cuevas, François Dupressoir, Luís Esquível,  
Ruben Gonzalez, **Benjamin Grégoire**, Andreas  
Hülsing, Vincent Hwang, Jan Jancar, Matthias  
Kannwischer, **Vincent Laporte**, **Jean-Christophe  
Léchenet**, Ting-han Lim, Cameron Low, Tiago  
Oliveira, Hugo Pacheco, Swarn Priya, Miguel  
Quaresma, Rolfe Schmidt, Antoine Séré, Lucas  
Tabary-Maujean, Pierre-Yves Strub, Yuval  
Yarom, Zhiyuan Zhang, **Jieyu Zheng**

≈9 papers, 30+ collaborators

Basavesh Ammanaghatta Shivakumar, **Santiago Arranz Olmos**, José Bacelar Almeida, Gustavo Xavier Delerue Marinho Alves, Manuel Barbosa, **Francisca Barros**, Gilles Barthe, Lionel Blatter, Chitchanok Chuengsatiansup, Ignacio Cuevas, François Dupressoir, Luís Esquível, Ruben Gonzalez, **Benjamin Grégoire**, Andreas Hülsing, Vincent Hwang, Jan Jancar, Matthias Kannwischer, **Vincent Laporte**, **Jean-Christophe Léchenet**, Ting-han Lim, Cameron Low, Tiago Oliveira, Hugo Pacheco, Swarn Priya, Miguel Quaresma, Rolfe Schmidt, Antoine Séré, Lucas Tabary-Maujean, Pierre-Yves Strub, Yuval Yarom, Zhiyuan Zhang, **Jieyu Zheng**

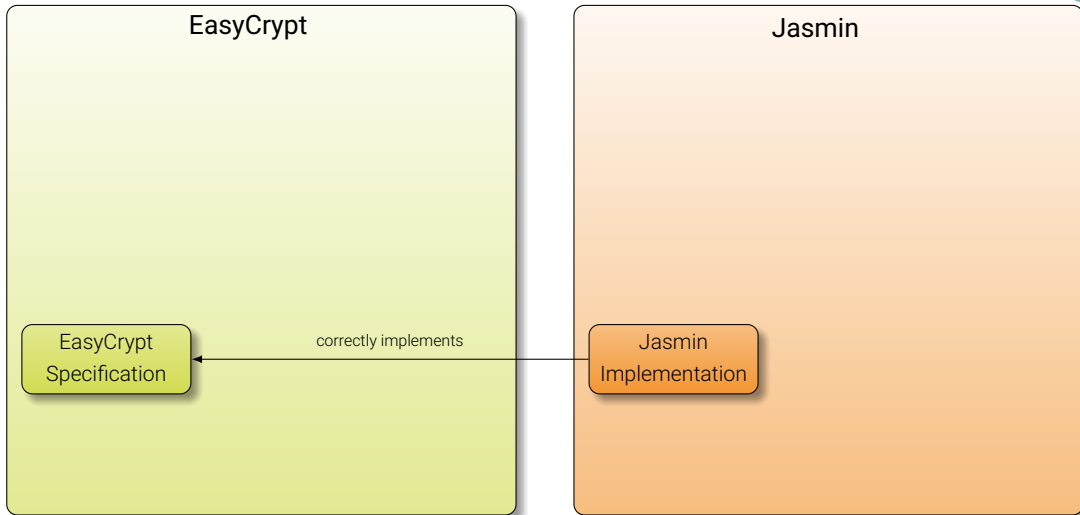


## FORMOSA CRYPTO



Universidade do Minho



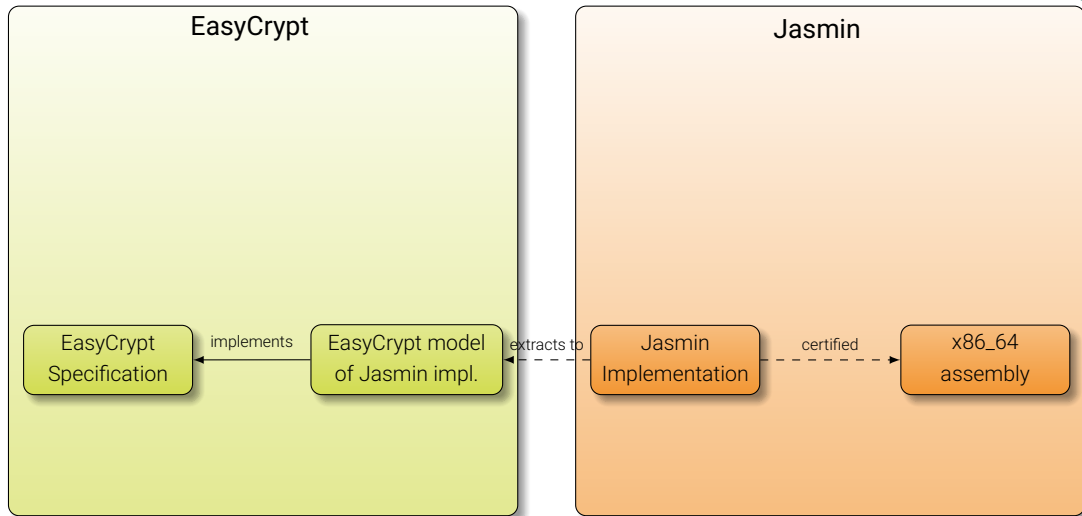




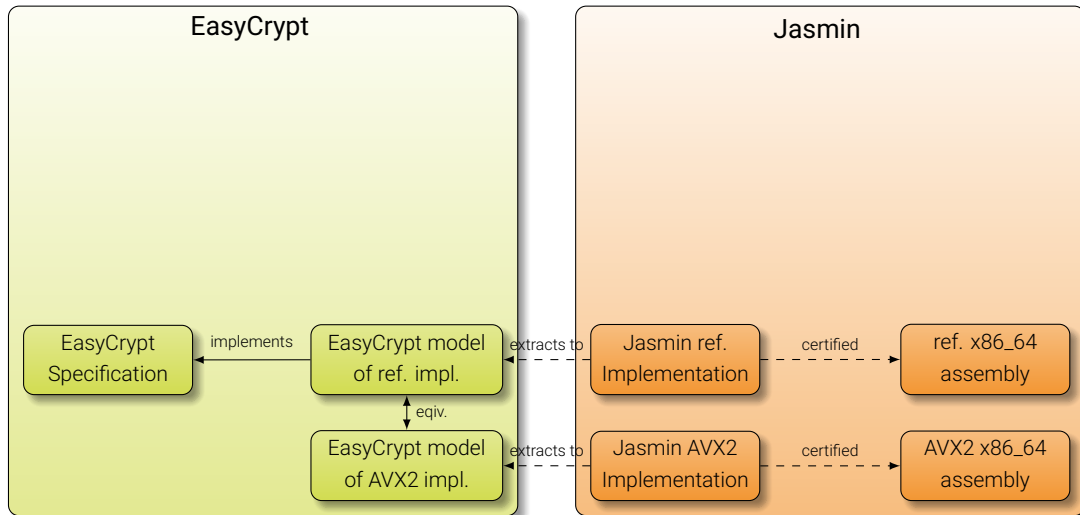
- ▶ Low-level programming language (“assembly in your head”)
- ▶ Predictable compilation
  - ▶ One line in Jasmin → one assembly instruction
  - ▶ Jasmin programs are architecture specific
  - ▶ Almost full control over generated assembly

- ▶ Low-level programming language (“assembly in your head”)
- ▶ Predictable compilation
  - ▶ One line in Jasmin → one assembly instruction
  - ▶ Jasmin programs are architecture specific
  - ▶ Almost full control over generated assembly
- ▶ High-level features: variables, functions, loops, conditionals. . .
- ▶ Register allocation (but no automatic spill-code)

- ▶ Low-level programming language (“assembly in your head”)
- ▶ Predictable compilation
  - ▶ One line in Jasmin → one assembly instruction
  - ▶ Jasmin programs are architecture specific
  - ▶ Almost full control over generated assembly
- ▶ High-level features: variables, functions, loops, conditionals. . .
- ▶ Register allocation (but no automatic spill-code)
- ▶ **Important for end-to-end guarantees:**
  - ▶ *Certified* compilation to assembly (proof in Rocq)
  - ▶ Extraction to EasyCrypt
  - ▶ Compilation preserves “constant time” (more later)



# Refining our high-level view



- ▶ Interactive verification framework
- ▶ Two languages, one functional, one imperative
- ▶ **Main purpose: Security reductions, game-hopping proofs**
  - ▶ Security goals and hardness assumptions as probabilistic programs
  - ▶ Support for standard arguments (e.g., hybrid arguments)
  - ▶ “Formally verify typical pen-and-paper security proofs”

- ▶ Interactive verification framework
- ▶ Two languages, one functional, one imperative
- ▶ **Main purpose: Security reductions, game-hopping proofs**
  - ▶ Security goals and hardness assumptions as probabilistic programs
  - ▶ Support for standard arguments (e.g., hybrid arguments)
  - ▶ “Formally verify typical pen-and-paper security proofs”
- ▶ **Additional feature: correctness proofs**
  - ▶ Relate functional and imperative programs
  - ▶ Prove equivalence of imperative programs

- ▶ Interactive verification framework
- ▶ Two languages, one functional, one imperative
- ▶ **Main purpose: Security reductions, game-hopping proofs**
  - ▶ Security goals and hardness assumptions as probabilistic programs
  - ▶ Support for standard arguments (e.g., hybrid arguments)
  - ▶ “Formally verify typical pen-and-paper security proofs”
- ▶ **Additional feature: correctness proofs**
  - ▶ Relate functional and imperative programs
  - ▶ Prove equivalence of imperative programs
- ▶ Proofs are manual, but with some automation (e.g., SMT-solver backend)





- ▶ Manuel translation from PDF (Kyber spec, FIPS 203) to EasyCrypt
- ▶ Uses both imperative and functional features
- ▶ Some interesting tradeoffs, e.g., input/output types for NTT
  - ▶ Strong typing: two different isomorphic rings
  - ▶ Weaker typing: both just coefficient arrays
  - ▶ Weaker typing allows in-place NTT



- ▶ Manuel translation from PDF (Kyber spec, FIPS 203) to EasyCrypt
- ▶ Uses both imperative and functional features
- ▶ Some interesting tradeoffs, e.g., input/output types for NTT
  - ▶ Strong typing: two different isomorphic rings
  - ▶ Weaker typing: both just coefficient arrays
  - ▶ Weaker typing allows in-place NTT
- ▶ Translation and checking could be avoided by **machine-readable standards**

# The specification



---

**Algorithm 14**  $\text{K-PKE.Encrypt}(\text{ek}_{\text{PKE}}, m, r)$ 

---

Uses the encryption key to encrypt a plaintext message using the randomness  $r$ .

**Input:** encryption key  $\text{ek}_{\text{PKE}} \in \mathbb{B}^{384k+32}$ .

**Input:** message  $m \in \mathbb{B}^{32}$ .

**Input:** randomness  $r \in \mathbb{B}^{32}$ .

**Output:** ciphertext  $c \in \mathbb{B}^{32(d_u k + d_v)}$ .

```
1:  $N \leftarrow 0$ 
2:  $\hat{\mathbf{t}} \leftarrow \text{ByteDecode}_{12}(\text{ek}_{\text{PKE}}[0 : 384k])$ 
3:  $\rho \leftarrow \text{ek}_{\text{PKE}}[384k : 384k + 32]$ 
4: for ( $i \leftarrow 0$ ;  $i < k$ ;  $i++$ )
5:   for ( $j \leftarrow 0$ ;  $j < k$ ;  $j++$ )
6:      $\mathbf{A}[i, j] \leftarrow \text{SampleNTT}(\rho \| j \| i)$ 
7:   end for
8: end for
9: for ( $i \leftarrow 0$ ;  $i < k$ ;  $i++$ )
10:   $\mathbf{y}[i] \leftarrow \text{SamplePolyCBD}_{\eta_1}(\text{PRF}_{\eta_1}(r, N))$ 
11:   $N \leftarrow N + 1$ 
12: end for
13: for ( $i \leftarrow 0$ ;  $i < k$ ;  $i++$ )
14:   $\mathbf{e}_1[i] \leftarrow \text{SamplePolyCBD}_{\eta_2}(\text{PRF}_{\eta_2}(r, N))$ 
15:   $N \leftarrow N + 1$ 
16: end for
17:  $\mathbf{e}_2 \leftarrow \text{SamplePolyCBD}_{\eta_2}(\text{PRF}_{\eta_2}(r, N))$ 
18:  $\hat{\mathbf{y}} \leftarrow \text{NTT}(\mathbf{y})$ 
19:  $\mathbf{u} \leftarrow \text{NTT}^{-1}(\hat{\mathbf{A}}^T \circ \hat{\mathbf{y}}) + \mathbf{e}_1$ 
20:  $\mu \leftarrow \text{Decompress}_1(\text{ByteDecode}_1(m))$ 
21:  $\mathbf{v} \leftarrow \text{NTT}^{-1}(\hat{\mathbf{t}}^T \circ \hat{\mathbf{y}}) + \mathbf{e}_2 + \mu$ 
22:  $c_1 \leftarrow \text{ByteEncode}_{d_u}(\text{Compress}_{d_u}(\mathbf{u}))$ 
23:  $c_2 \leftarrow \text{ByteEncode}_{d_v}(\text{Compress}_{d_v}(\mathbf{v}))$ 
24: return  $c \leftarrow (c_1 \| c_2)$ 
```

---

```
proc enc_derand(pk : pkey, m : plaintext, r : W8.t Array32.t) : ciphertext = {
  (tv, rho)  $\leftarrow$  pk;
  _N  $\leftarrow$  0;
  that1  $\leftarrow$  EncDec.decode12_vec(tv);
  that  $\leftarrow$  ofipolyvec that1;
  i  $\leftarrow$  0;
  while (i < kvec) {
    j  $\leftarrow$  0;
    while (j < kvec) {
      XOF(O).init(rho, W8.of_int i, W8.of_int j);
      c  $\leftarrow$  Parse(XOF(O).sample());
      aT[[i, j]]  $\leftarrow$  c;
      j  $\leftarrow$  j + 1;
    }
    i  $\leftarrow$  i + 1;
  }
  i  $\leftarrow$  0;
  while (i < kvec) {
    c  $\leftarrow$  CBD2(PRF).sample(r, _N);
    yv  $\leftarrow$  set yv i c;
    _N  $\leftarrow$  _N + 1;
    i  $\leftarrow$  i + 1;
  }
  i  $\leftarrow$  0;
  while (i < kvec) {
    c  $\leftarrow$  CBD2(PRF).sample(r, _N);
    e1  $\leftarrow$  set e1 i c;
    _N  $\leftarrow$  _N + 1;
    i  $\leftarrow$  i + 1;
  }
  e2  $\leftarrow$  CBD2(PRF).sample(r, _N);
  yhat  $\leftarrow$  nttv yv;
  u  $\leftarrow$  invnttv (ntt_mm mul aT yhat) + e1;
  mp  $\leftarrow$  EncDec.decode1(m);
  v  $\leftarrow$  invntt (ntt_dotp that yhat) + e2 + decompress_poly 1 mp;
  c1  $\leftarrow$  EncDec.encode10_vec(compress_polyvec 10 u);
  c2  $\leftarrow$  EncDec.encode4(compress_poly 4 v);
  return (c1, c2);
}
```



Goal: Link specification to IND-CCA notion

## Goal: Link specification to IND-CCA notion

### Step 1: IND-CPA security of K-PKE

- ▶ Typical game-hopping proof for abstract LWE-based PKE
- ▶ Refine to concrete parameters of K-PKE
- ▶ Reduction from *hashed MLWE*

## Goal: Link specification to IND-CCA notion

### Step 1: IND-CPA security of K-PKE

- ▶ Typical game-hopping proof for abstract LWE-based PKE
- ▶ Refine to concrete parameters of K-PKE
- ▶ Reduction from *hashed MLWE*

### Step 2: Decryption-failure bounds for K-PKE

- ▶ Original analysis assumes compression of uniform coefficients
- ▶ Cannot rely on MLWE assumption in correctness game
- ▶ Move from adversarial notion to statistic notion
- ▶ Distribution too complex to fully compute

## Goal: Link specification to IND-CCA notion

### Step 1: IND-CPA security of K-PKE

- ▶ Typical game-hopping proof for abstract LWE-based PKE
- ▶ Refine to concrete parameters of K-PKE
- ▶ Reduction from *hashed MLWE*

### Step 2: Decryption-failure bounds for K-PKE

- ▶ Original analysis assumes compression of uniform coefficients
- ▶ Cannot rely on MLWE assumption in correctness game
- ▶ Move from adversarial notion to statistic notion
- ▶ Distribution too complex to fully compute

### Step 3: IND-CCA security of ML-KEM

- ▶ Prove IND-CCA for MLKEM\_OP
- ▶ Central difference: SHA3-512 replaced by RO

## Spec $\leftrightarrow$ ref impl.

- ▶ Montgomery and Barrett reductions
- ▶ Lazy reductions (bounds checking!)
- ▶ “Funny” code, e.g., for  $a = a/3329$

```
a *= 80635;
```

```
a >>= 28;
```

- ▶ Magic values in NTT code
- ▶ Link functional and imperative code



## Spec $\leftrightarrow$ ref impl.

- ▶ Montgomery and Barrett reductions
- ▶ Lazy reductions (bounds checking!)
- ▶ “Funny” code, e.g., for  $a = a/3329$

```
a *= 80635;
```

```
a >>= 28;
```

- ▶ Magic values in NTT code
- ▶ Link functional and imperative code

## ref. impl. $\leftrightarrow$ AVX2 impl.

- ▶ Optimizations *across* functions
- ▶ Different internal data representations
- ▶ Non-word-size operations (compress)
- ▶ Complex vectorized rejection sampling

## Spec $\leftrightarrow$ ref impl.

- ▶ Montgomery and Barrett reductions
- ▶ Lazy reductions (bounds checking!)
- ▶ “Funny” code, e.g., for  $a = a/3329$

```
a *= 80635;
```

```
a >>= 28;
```

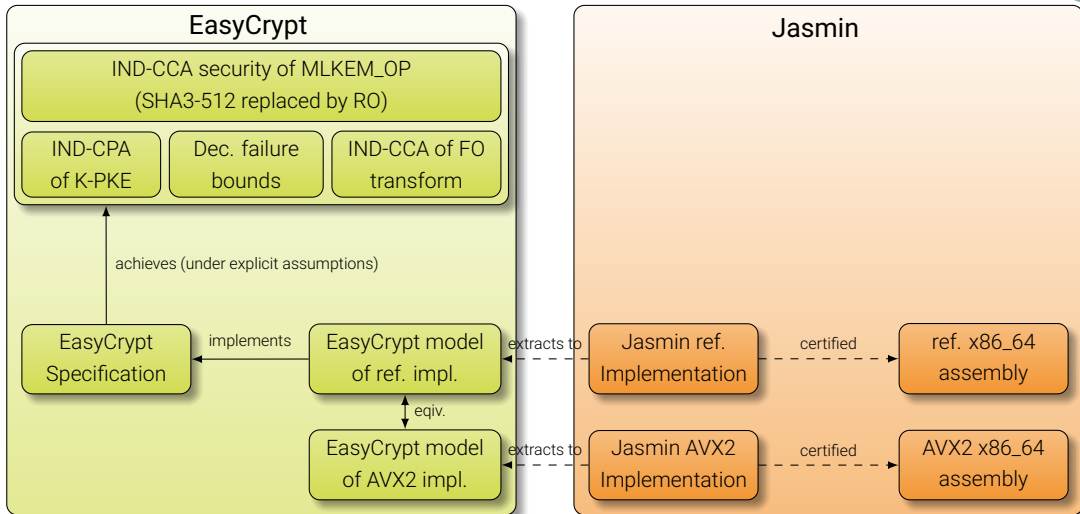
- ▶ Magic values in NTT code
- ▶ Link functional and imperative code

## ref. impl. $\leftrightarrow$ AVX2 impl.

- ▶ Optimizations *across* functions
- ▶ Different internal data representations
- ▶ Non-word-size operations (compress)
- ▶ Complex vectorized rejection sampling

- ▶ Massive manual effort with quite some student frustration
- ▶ Significant improvement 2025:
  - ▶ *automated* circuit-equivalence checking
  - ▶ integrated in EasyCrypt, interleave with deductive reasoning
  - ▶ faster verification, more robust proofs, easier proof maintenance

# So, where are we?



## The old way

- ▶ Static safety check:  
`jasminc -checksafety`
- ▶ Great when it works
- ▶ Takes a long time (not modular)
- ▶ Often fails for safe code

## The old way

- ▶ Static safety check:  
`jasminc -checksafety`
- ▶ Great when it works
- ▶ Takes a long time (not modular)
- ▶ Often fails for safe code

## The new way

- ▶ Master's thesis by **Francisca Barros**
- ▶ Modular design
- ▶ Safety contracts and assertions in Jasmin
- ▶ Automatic discharge of assertions
- ▶ Prove remaining assertions in EasyCrypt

```
fn _gen_matrix_avx2
( reg mut ptr u16[MLKEM_K * MLKEM_K * MLKEM_N] matrix
, reg const ptr u8[32] rho
, #spill_to_mmx reg u64 transposed
) -> reg ptr u16[MLKEM_K * MLKEM_K * MLKEM_N]
requires {is_arr_init(rho,0,32) && 0<= transposed && transposed <= 1}
ensures {is_arr_init(result.0,0,MLKEM_K * MLKEM_K * MLKEM_N * 2)}
{
    ...
}
```

- ▶ Enforce constant-time on Jasmin source level
- ▶ Every piece of data is either **secret** or **public**
- ▶ Flow of secret information is traced by type system

*“Any operation with a secret input produces a secret output”*

- ▶ Enforce constant-time on Jasmin source level
- ▶ Every piece of data is either `secret` or `public`
- ▶ Flow of secret information is traced by type system
  - “Any operation with a secret input produces a secret output”*
- ▶ Branch conditions and memory indices need to be `public`

- ▶ Enforce constant-time on Jasmin source level
- ▶ Every piece of data is either `secret` or `public`
- ▶ Flow of secret information is traced by type system
  - “Any operation with a secret input produces a secret output”*
- ▶ Branch conditions and memory indices need to be `public`
- ▶ In principle can do this also in, e.g., Rust (`secret_integers` crate)
- ▶ **Jasmin compiler preserves constant-time!**



- ▶ Enforce constant-time on Jasmin source level
- ▶ Every piece of data is either `secret` or `public`
- ▶ Flow of secret information is traced by type system
  - “Any operation with a secret input produces a secret output”*
- ▶ Branch conditions and memory indices need to be `public`
- ▶ In principle can do this also in, e.g., Rust (`secret_integers` crate)
- ▶ **Jasmin compiler preserves constant-time!**
- ▶ Explicit `#declassify` primitive to move from `secret` to `public`
- ▶ `#declassify` creates a proof obligation!

- ▶ Some arithmetic instructions also leak through timing
- ▶ Example: `DIV` (exploited in “KyberSlash”, CHES 2025)
- ▶ Which instructions are safe to use, *also for future CPUs*?

- ▶ Some arithmetic instructions also leak through timing
- ▶ Example: `DIV` (exploited in “KyberSlash”, CHES 2025)
- ▶ Which instructions are safe to use, *also for future CPUs*?
- ▶ Intel Core Security Team, 2022: Data Operand Independent Timing (DOIT)
  - ▶ Subset of x86\_64 instructions **guaranteed** to not leak through timing
  - ▶ Requires switching CPU to DOIT mode

- ▶ Some arithmetic instructions also leak through timing
- ▶ Example: `DIV` (exploited in “KyberSlash”, CHES 2025)
- ▶ Which instructions are safe to use, *also for future CPUs*?
- ▶ Intel Core Security Team, 2022: Data Operand Independent Timing (DOIT)
  - ▶ Subset of x86\_64 instructions **guaranteed** to not leak through timing
  - ▶ Requires switching CPU to DOIT mode
- ▶ Support for DOIT in Jasmin compiler
  - ▶ Use only DOIT instructions on secret inputs
  - ▶ No serious limitation for most *optimized* code
  - ▶ All relevant vector instructions are DOIT

```
stack u8[10] public;
stack u8[32] secret;
reg u8 t;
reg u64 r, i;

i = 0;
while(i < 10) {
    t = public[(int) i] ;
    r = leak(t);
    ...
}
```

## Fencing

- ▶ Can prevent speculation through **barriers** (LFENCE)
- ▶ Protecting *all* branches is possible but costly

## Fencing

- ▶ Can prevent speculation through **barriers** (LFENCE)
- ▶ Protecting *all* branches is possible but costly

## Speculative Load Hardening

- ▶ Idea: maintain misprediction predicate `ms` (in a register)
- ▶ At every branch use arithmetic to update predicate
- ▶ Option 1: Mask every loaded value with `ms`
- ▶ Option 2: Mask every address with `ms`
- ▶ Effect: during misspeculation “leak” constant value

## Fencing

- ▶ Can prevent speculation through **barriers** (LFENCE)
- ▶ Protecting *all* branches is possible but costly

## Speculative Load Hardening

- ▶ Idea: maintain misprediction predicate `ms` (in a register)
- ▶ At every branch use arithmetic to update predicate
- ▶ Option 1: Mask every loaded value with `ms`
- ▶ Option 2: Mask every address with `ms`
- ▶ Effect: during misspeculation “leak” constant value
- ▶ Implemented in LLVM since version 8
  - ▶ Still large performance overhead
  - ▶ No formal guarantees of security



Do we need to mask/protect all loads?

**Do we need to mask/protect all loads?**

- ▶ No need to mask loads into registers that never enter leaking instructions

## Do we need to mask/protect all loads?

- ▶ No need to mask loads into registers that never enter leaking instructions
- ▶ `secret` registers never enter leaking instructions!
- ▶ Obvious idea: mask only loads into `public` registers

- ▶ Type system gets three security levels:
  - ▶ **secret**: secret
  - ▶ **public**: public, also during misspeculation
  - ▶ **transient**: public, but possibly secret during misspeculation

- ▶ Type system gets three security levels:
  - ▶ `secret`: secret
  - ▶ `public`: public, also during misspeculation
  - ▶ `transient`: public, but possibly secret during misspeculation
- ▶ Maintain misspeculation flag `ms`:
  - ▶ `ms = #init_msf()`: Translate to `LFENCE`, set `ms` to zero
  - ▶ `ms = #set_msf(b, ms)`: Set `ms` according to branch condition `b`
  - ▶ Branches invalidate `ms`

- ▶ Type system gets three security levels:
  - ▶ `secret`: secret
  - ▶ `public`: public, also during misspeculation
  - ▶ `transient`: public, but possibly secret during misspeculation
- ▶ Maintain misspeculation flag `ms`:
  - ▶ `ms = #init_msfn()`: Translate to `LFENCE`, set `ms` to zero
  - ▶ `ms = #set_msfn(b, ms)`: Set `ms` according to branch condition `b`
  - ▶ Branches invalidate `ms`
- ▶ Two operations to lower level:
  - ▶ `x = #protect(x, ms)`: Go from `transient` to `public`
  - ▶ `#protect` translates to mask by `ms`

- ▶ Type system gets three security levels:
  - ▶ `secret`: secret
  - ▶ `public`: public, also during misspeculation
  - ▶ `transient`: public, but possibly secret during misspeculation
- ▶ Maintain misspeculation flag `ms`:
  - ▶ `ms = #init_msf()`: Translate to `LFENCE`, set `ms` to zero
  - ▶ `ms = #set_msf(b, ms)`: Set `ms` according to branch condition `b`
  - ▶ Branches invalidate `ms`
- ▶ Two operations to lower level:
  - ▶ `x = #protect(x, ms)`: Go from `transient` to `public`
  - ▶ `#protect` translates to mask by `ms`
  - ▶ `#declassify r`: Go from `secret` to `transient`
  - ▶ `#declassify` requires cryptographic proof/argument

- ▶ Type system gets three security levels:
  - ▶ `secret`: secret
  - ▶ `public`: public, also during misspeculation
  - ▶ `transient`: public, but possibly secret during misspeculation
- ▶ Maintain misspeculation flag `ms`:
  - ▶ `ms = #init_msf()`: Translate to `LFENCE`, set `ms` to zero
  - ▶ `ms = #set_msf(b, ms)`: Set `ms` according to branch condition `b`
  - ▶ Branches invalidate `ms`
- ▶ Two operations to lower level:
  - ▶ `x = #protect(x, ms)`: Go from `transient` to `public`
  - ▶ `#protect` translates to mask by `ms`
  - ▶ `#declassify r`: Go from `secret` to `transient`
  - ▶ `#declassify` requires cryptographic proof/argument
- ▶ Still: allow branches and indexing only for `public`
- ▶ Also, allow non-DOIT instructions only on `public` inputs



- ▶ We know what inputs are **secret** and what inputs are **public**
- ▶ Most of the state is actually **secret**
- ▶ Most loads do not need **protect**!

- ▶ We know what inputs are **secret** and what inputs are **public**
- ▶ Most of the state is actually **secret**
- ▶ Most loads do not need **protect**!
- ▶ Even better: mark additional inputs as **secret**
- ▶ No cost if those inputs don't flow into leaking instructions

- ▶ We know what inputs are **secret** and what inputs are **public**
- ▶ Most of the state is actually **secret**
- ▶ Most loads do not need **protect**!
- ▶ Even better: mark additional inputs as **secret**
- ▶ No cost if those inputs don't flow into leaking instructions
- ▶ Even better: Spills don't need **protect** if there is no branch between store and load

- ▶ We know what inputs are **secret** and what inputs are **public**
- ▶ Most of the state is actually **secret**
- ▶ Most loads do not need **protect**!
- ▶ Even better: mark additional inputs as **secret**
- ▶ No cost if those inputs don't flow into leaking instructions
- ▶ Even better: Spills don't need **protect** if there is no branch between store and load
- ▶ Even better: "Spill" public data to MMX registers instead of stack

# If there's a Spectre v1...



- ▶ Spectre v2: indirect branches (not supported by Jasmin)
- ▶ Spectre v3, aka “Meltdown”: fix in HW and firmware
- ▶ Spectre v4: speculative store bypass (disable with SSBD)

# If there's a Spectre v1...



- ▶ Spectre v2: indirect branches (not supported by Jasmin)
- ▶ Spectre v3, aka “Meltdown”: fix in HW and firmware
- ▶ Spectre v4: speculative store bypass (disable with SSBD)
- ▶ **Spectre-RSB**: function return speculates *anywhere*
- ▶ Attacker can choose to speculate right behind any defense!

- ▶ Spectre v2: indirect branches (not supported by Jasmin)
- ▶ Spectre v3, aka “Meltdown”: fix in HW and firmware
- ▶ Spectre v4: speculative store bypass (disable with SSBD)
- ▶ **Spectre-RSB**: function return speculates *anywhere*
- ▶ Attacker can choose to speculate right behind any defense!
- ▶ Solution in Jasmin:
  - ▶ Rewrite all returns through “branch table”
  - ▶ Implement branch table through *conditional* branches
  - ▶ Speculate only to one of the call sites
  - ▶ At call sites all **public** values become **transient**
  - ▶ Use **#protect** and **ms**



*“...A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys”*

—FIPS 140-3, Section 9.7.A



*"...A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys"*

—FIPS 140-3, Section 9.7.A

## Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

*"...A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys"*

—FIPS 140-3, Section 9.7.A

## Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

## Failure modes

0. Don't perform any zeroization

*"...A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys"*

—FIPS 140-3, Section 9.7.A

## Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

## Failure modes

0. Don't perform any zeroization
1. Dead-store elimination

*"...A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys"*

—FIPS 140-3, Section 9.7.A

## Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

## Failure modes

0. Don't perform any zeroization
1. Dead-store elimination
2. Only API-level stack zeroization

*"...A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys"*

—FIPS 140-3, Section 9.7.A

## Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

## Failure modes

0. Don't perform any zeroization
1. Dead-store elimination
2. Only API-level stack zeroization
3. Don't scrub source-level invisible data

*"...A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys"*

—FIPS 140-3, Section 9.7.A

## Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

## Failure modes

0. Don't perform any zeroization
1. Dead-store elimination
2. Only API-level stack zeroization
3. Don't scrub source-level invisible data
4. Mis-estimate stack space when scrubbing from caller

## Solution in Jasmin compiler

Zeroize used stack space and registers when returning from export function

## Solution in Jasmin compiler

**Zeroize used stack space and registers when returning from export function**

- ▶ Make use of multiple features of Jasmin:
  - ▶ Compiler has global view
  - ▶ All stack usage is known at compile time
  - ▶ Entry/return point is clearly defined

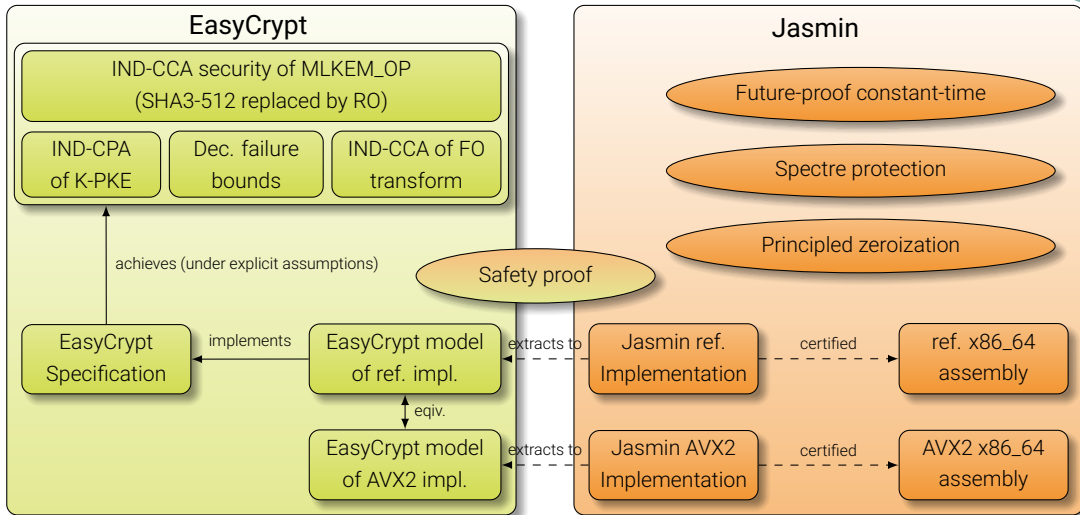


## Solution in Jasmin compiler

Zeroize used stack space and registers when returning from export function

- ▶ Make use of multiple features of Jasmin:
  - ▶ Compiler has global view
  - ▶ All stack usage is known at compile time
  - ▶ Entry/return point is clearly defined
- ▶ Performance overhead for Kyber768 (on Comet Lake):
  - ▶ 0.59% for Keypair
  - ▶ 0.24% for Encaps
  - ▶ 1.04% for Decaps

# The next generation of crypto software



## Cycles for ML-KEM-768

CPU	Implementation	keypair	enc	dec
8700K	Jasmin AVX2*	40134	40599	43437
	pq-crystals	39722	39761	46161
11700K	Jasmin AVX2*	37458	37798	39970
	pq-crystals	36958	38082	42566
13900K	Jasmin AVX2*	34732	35212	43784
	pq-crystals	31448	32090	36064

\*with Spectre-v1 (without Spectre-RSB) protections

## Work in progress (selection)

- ▶ Deployment in Signal's contact discovery (RWC talk upcoming!)
- ▶ Integrate Spectre-RSB protections
- ▶ Implement in **crypto agent** process
- ▶ Extend to more architectures and more primitives (next up: ML-DSA)
- ▶ Interface to super-optimizers (e.g., CryptOpt, SLOTHY)
- ▶ Improve usability/scalability of tools


## Longer-term plans (selection)

- ▶ Guarantee preservation of speculative constant time
- ▶ Masked implementations
- ▶ Machine-readable standards (?)

- ▶ Barbosa, Kannwischer, Lim, Schwabe, Strub. *Formally Verified Correctness Bounds for Lattice-Based Cryptography*. ACM CCS 2025.
- ▶ Almeida, Marinho Alves, Barbosa, Barthe, Esquível, Hwang, Oliveira, Pacheco, Schwabe, Strub. *Faster Verification of Faster Implementations: Combining Deductive and Circuit-Based Reasoning in EasyCrypt*. IEEE S&P 2025.
- ▶ Arranz Olmos, Barthe, Grégoire, Jancar, Laporte, Oliveira, Schwabe. *Let's DOIT: Using Intel's Extended HW/SW Contract for Secure Compilation of Crypto Code*. TCHES 2025-3.
- ▶ Arranz Olmos, Barthe, Chuengsatiansup, Grégoire, Laporte, Oliveira, Schwabe, Yarom, Zhang. *Protecting Cryptographic Code Against Spectre-RSB (and, in Fact, All Known Spectre Variants)*. ASPLOS 2025.
- ▶ Barbosa, Schwabe. *Kyber terminates*. Polynesian Journal of Mathematics, vol. 1, issue 6 (2024).

- ▶ Almeida, Arranz Olmos, Barbosa, Barthe, Dupressoir, Grégoire, Laporte, Léchenet, Low, Oliveira, Pacheco, Quaresma, Schwabe, Strub. *Formally verifying Kyber – Episode V: Machine-checked IND-CCA security and correctness of ML-KEM in EasyCrypt*. CRYPTO 2024.
- ▶ Arranz Olmos, Barthe, Gonzalez, Grégoire, Laporte, Léchenet, Oliveira, Schwabe. *High-assurance zeroization*. TCHES 2024-1.
- ▶ Almeida, Barbosa, Barthe, Grégoire, Laporte, Léchenet, Oliveira, Pacheco, Quaresma, Schwabe, Séré, Strub. *Formally verifying Kyber – Episode IV: Implementation Correctness*. TCHES 2023-3.
- ▶ Ammanaghata Shivakumar, Barthe, Grégoire, Laporte, Oliveira, Priya, Schwabe, Tabary-Maujean. *Typing High-Speed Cryptography against Spectre v1*. IEEE S&P 2023.

# Thank you!



[https://github.com/pq-code-package/  
mlkem-libjade](https://github.com/pq-code-package/mlkem-libjade)



<https://formosa-crypto.org>